

An Axiomatic Approach to Fairness

Tian Lan¹, David Kao², Mung Chiang¹, Ashutosh Sabharwal²

¹Department of Electrical Engineering, Princeton University, NJ 08544, USA

²Department of Electrical Engineering, Rice University, TX 77005, USA

ABSTRACT

We present a set of five axioms for fairness measures in resource allocation. A family of fairness measures satisfying the axioms is constructed. Well-known notions such as α -fairness, Jain's index and entropy are shown to be special cases. Properties of fairness measures satisfying the axioms are proven, including Schur-concavity. Among the engineering implications is a new understanding of α -fair utility functions and an interpretation of "larger α is more fair".

1. QUANTIFYING FAIRNESS

Given a vector $\mathbf{x} \in \mathcal{R}_+^n$, where x_i is the resource allocated to user i , *how* fair is it?

Given a vector $\mathbf{x} \in \mathbb{R}_+^n$, where x_i is the resource allocated to user i , *how* fair is it?

One approach to quantify a degree of fairness associated with \mathbf{x} is through a fairness measure, which is a function f that maps \mathbf{x} into a real number. Various fairness measures have been proposed throughout the years, e.g., in [1, 2, 3, 4, 5, 6]. These range from simple ones, e.g., the ratio between the smallest and the largest entries of \mathbf{x} , to more sophisticated functions, e.g., Jain's index and entropy function. Some of these fairness measures map \mathbf{x} to normalized numbers between 0 and 1, where 0 denotes the minimum fairness, 1 denotes the maximum fairness, often corresponding to an \mathbf{x} where all x_i are the same, and a larger value indicate more fairness. For example, min-max ratio [1] is given by the maximum ratio of any two user's resource allocation, while Jain's index [3] computes a normalized square mean. Are these fairness measures related to each other? Is one of them measures "better" than the other? What other measures of fairness may be useful?

Another approach that has gained attention in the networking research community since [7, 8] is α -fairness and the associated utility maximization. A maximizer of α -fair utility function satisfies the definition of α -fairness. Two well-known examples are as follows: a maximizer of the log utility function ($\alpha = 1$) is proportional fair (i.e., any change in \mathbf{x} has a negative total normalized change), and a maxi-

mizer of the α -fair utility function with $\alpha \rightarrow \infty$ is max-min fair (i.e., it is impossible to increase an x_i without decreasing some $x_j \leq x_i$). More recently, α -fair utility functions have also been connected to divergence measures [9]. As in [10, 11], parameter α can be viewed as a fairness measure in the sense that a fairer allocation is one with a larger α , although the exact role of α in trading-off fairness and throughput can sometimes be surprising [12]. While it is often held as self-evident that $\alpha \rightarrow \infty$ is more fair than $\alpha = 1$, which is in turn more fair than $\alpha = 0$, it remains unclear what does it mean to say that $\alpha = 3$ is more fair than $\alpha = 2$?

These two approaches for quantifying fairness are obviously different. Since α -fair utility functions are continuous and strictly increasing in each entry of \mathbf{x} , maximizing them results in Pareto optimal resource allocations. On the other hand, fairness measures that maps \mathbf{x} to a normalized range are independent to the absolute values of \mathbf{x} . What notions of fairness or efficiency is captured by each approach? Can the two approaches be unified?

To address the above questions, we develop an axiomatic approach to measure fairness. We discover that a set of five axioms, each of which simple and intuitive, can lead to a useful family of fairness measures. The axioms are: the axiom of continuity, of homogeneity, of asymptotic saturation, of irrelevance of partition, and of monotonicity. Starting with these five self-evident truths, we can *generate* a family of fairness measures from a generator function g : any increasing and continuous functions that lead to a well-defined "mean" function (i.e., from any Kolmogorov-Nagumo function [16]). For example, using power functions with exponent β as the generator function, we derive a unique family of fairness measures f_β that include all of the following as special cases, depending on the choice of β : Jain's index, maximum or minimum ratio, entropy, and α -fair utility, and reveals new fairness measures corresponding to other ranges of β .

We note that our approach, unlike other well known axiomatic constructions such as the Nash bargaining solution [18] and Shapley value [19], specifies a broad class of fairness measures rather than an optimality concept. This is a direct consequence of one axiom which by construction removes the concept of efficiency from fairness. However, we demonstrate that in many optimization theoretic constructions, a fairness measure is incorporated into the objective function. For example, we show that an α -fair utility function (which includes proportional fairness [7] and thus the Nash Bargaining Solution as a special case) can be factorized as the product of two components: our fairness measure with $\alpha = \beta$ and a function of the total throughput. We also

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

show that such a factorization indicates a tradeoff between fairness and efficiency, with α -fairness achieving Pareto dominance with the maximum possible fairness. This facilitates a clearer understanding of what it means to say that a larger α is “more fair” for general $\alpha \in [0, \infty)$.

The axiomatic construction of fairness measures also illuminates their engineering implications. Any fairness measure satisfying the five axioms can be proven to have many useful properties, including Schur-concavity [14]. Consequently, any operation balancing resources between two user always results in a higher fairness value, extending previous results using majorization to characterize fairness [4, 13].

Main notation is shown in Table 1. Proofs can be found in the technical report [21] available online.

Variable	Meaning
\mathbf{x}	Resource allocation vector of length n
\mathbf{x}^\uparrow	Sorted vector with smallest element being first
$w(\mathbf{x})$	Sum of all elements of \mathbf{x}
$f(\cdot), f_\beta(\cdot)$	Fairness measure (of parameter β)
$g(\cdot)$	Generator function
s_i	Positive weights for weighted mean
$\mathbf{1}_n$	Vector of all ones of length n
$\mathbf{x} \succeq \mathbf{y}$	Vector \mathbf{x} majorizes vector \mathbf{y}
β	Parameter for power function $g(y) = y^\beta$
$U_\alpha(\cdot)$	α -fair utility with parameter α
$H(\cdot)$	Shannon entropy function
$J(\cdot)$	Jain’s index
$\Phi(\cdot)$	Our utility for fairness and efficiency

Table 1: Table of notations in this paper.

2. AXIOMS

Let \mathbf{x} be a resource allocation vector with n non-negative elements. A fairness measure $f(\mathbf{x})$ is a mapping from \mathbf{x} to a real number, i.e., $f : \mathbb{R}_+^n \rightarrow \mathbb{R}$, for all integer $n \geq 1$. We first introduce the following set of fairness axioms, whose explanations and implications are given next.

- 1) *Axiom of Continuity.* Fairness measure $f(\mathbf{x})$ is continuous on \mathbb{R}_+^n for all integer $n \geq 1$.
- 2) *Axiom of Homogeneity.*

$$f(\mathbf{x}) = f(t \cdot \mathbf{x}), \quad \forall t > 0 \quad (1)$$

Without loss of generality, we have $|f(\mathbf{1})| = 1$ **Tian: is the first 1 a vector?** since fairness is a constant for $n = 1$.

- 3) *Axiom of Asymptotic Saturation.*

$$\lim_{n \rightarrow \infty} \frac{f(\mathbf{1}_{n+1})}{f(\mathbf{1}_n)} = 1. \quad (2)$$

- 4) *Axiom of Irrelevance of Partition.* If we divide \mathbf{x} into two parts $\mathbf{x} = [\mathbf{x}^1, \mathbf{x}^2]$, fairness $f(\mathbf{x}^1, \mathbf{x}^2)$ can be computed recursively (with respect to a generator function $g(y)$) and is independent of the partition, i.e.,

$$f(\mathbf{x}^1, \mathbf{x}^2) = f(w(\mathbf{x}^1), w(\mathbf{x}^2)) \cdot g^{-1} \left(\sum_{i=1}^2 s_i \cdot g(f(\mathbf{x}^i)) \right), \quad (3)$$

where $w(\mathbf{x}^1)$ and $w(\mathbf{x}^2)$ denote the sum of vector \mathbf{x}^1 and \mathbf{x}^2 , respectively, $g(y)$ is continuous and strictly monotonic function that can generate the following function h :

$$h = g^{-1} \left(\sum_{i=1}^2 s_i \cdot g(f(\mathbf{x}^i)) \right), \quad (4)$$

with positive weights satisfying $\sum_i s_i = 1$ such that h qualifies as a *mean* function [15] of $\{f(\mathbf{x}^i), \forall i\}$.

- 5) *Axiom of Monotonicity.* For $n = 2$ users, fairness measure $f(\theta, 1 - \theta)$ is monotonically increasing for $\theta \in [0, \frac{1}{2}]$.

Axioms 1-2 are very intuitive. The Axiom of Continuity says that a slight change in resource allocation show up in the fairness measure. The Axiom of Homogeneity says that the fairness measure is independent of the unit of measurement.

Axiom 3 is used to ensure uniqueness. Suppose $f(\mathbf{x})$ is a fairness measure with an exponential scaling behavior and satisfies all other axioms (with respect to a mean function $g(y)$) except for Axiom 3. It is easy to see that by making logarithmic change of variable, fairness measure $\log f(\mathbf{x})$ satisfies all axioms above with respect to a mean function $e^{g(y)}$. Therefore, Axiom 3 picks up fairness measures with a particular scaling order and allows invariance under change of variable.

Axiom 4 is the one concerning the scaling of fairness. A natural choice of the weight s_i in (3) is to choose the value proportional to the sum resource of vector \mathbf{x}^i . In general, we consider the following weights

$$s_i = \frac{w^\rho(\mathbf{x}^i)}{\sum_j w^\rho(\mathbf{x}^j)}, \quad \forall i \quad (5)$$

where ρ is an arbitrary exponent. When $\rho = 0$, weights in (5) are equal and lead to un-weighted mean in Axiom 4. As shown in Section 4, the parameter ρ can be chosen such that the hierarchical computation is independent of partition as stated in Axiom 4.

Axiom 5 is the only axiom that actually involves a value statement: when there are just two users, more equalized is more fair.

By definition, axioms are true. However, not all sets of axioms are useful: recovering all known notions and discovering new measures and properties. We start showing the above six axioms are useful with the following existence and uniqueness results.

THEOREM 1. (Existence.) *There exists a fairness measure $f(\mathbf{x})$ satisfying Axioms 1-5. Furthermore, the fairness*

achieved by equal-resource allocations $\mathbf{1}_n$ is independent of the choice of $g(y)$, i.e.,

$$f(\mathbf{1}_n) = n^r \cdot f(\mathbf{1}), \quad \forall n \geq 1, \quad (6)$$

where r is a constant exponent.

THEOREM 2. (Uniqueness.) Given a generator function g , the resulting $f(\mathbf{x})$ satisfying Axioms 1-5 is unique.

3. PROPERTIES OF FAIRNESS MEASURES

We first prove an intuitive corollary from the axioms that will be useful for the rest of the presentation.

COROLLARY 1. (Symmetry.) A fairness measure satisfying Axioms 1-5 is symmetric over \mathbf{x} :

$$f(x_1, x_2, \dots, x_n) = f(x_{i_1}, x_{i_2}, \dots, x_{i_n}), \quad (7)$$

where i_1, \dots, i_n is an arbitrary permutation of indices $1, \dots, n$.

The symmetry property shows that the fairness measure $f(\mathbf{x})$ satisfying Axioms 1-5 is irrelevant of labeling of users. Therefore, for fixed population size n , the fairness of two resource allocation vectors \mathbf{x} and \mathbf{y} are compared through the distribution of their elements.

Majorization [14] is a partial order over vectors. The main purpose was to study the general notion of whether the elements of vector \mathbf{x} are less spread out than the elements of vector \mathbf{y} . We say that \mathbf{x} is majorized by \mathbf{y} , and we write $\mathbf{x} \preceq \mathbf{y}$, if $\sum_{i=1}^n x_i = \sum_{i=1}^n y_i$ (always satisfied due to Axiom 2) and

$$\sum_{i=1}^d x_i^\uparrow \leq \sum_{i=1}^d y_i^\uparrow, \quad \text{for } d = 1, \dots, n, \quad (8)$$

where x_i^\uparrow and y_i^\uparrow are the i th elements of \mathbf{x}^\uparrow and \mathbf{y}^\uparrow , sorted in ascending order. According to this definition, among the vectors with the same sum of elements, one with the equal elements is the mostly majorized vector.

Intuitively, $\mathbf{x} \preceq \mathbf{y}$ can be interpreted as \mathbf{y} being a fairer allocation than \mathbf{x} . It is a classical result [14] that \mathbf{x} is majorized by \mathbf{y} , if and only if, from \mathbf{x} we can produce \mathbf{y} by a finite sequence of *Robin Hood operations*, where we replace two elements x_i and $x_j < x_i$ with $x_i - \epsilon$ and $x_j + \epsilon$, respectively, for some $\epsilon \in (0, x_i - x_j)$. Majorization cannot be used to define a reliable fairness index (due to the fact that it is a partial order and fails to compare vectors in certain cases). Still, if resource allocations \mathbf{x} is majorized by \mathbf{y} , it is desirable to have a fairness measure $f(\mathbf{x}) \leq f(\mathbf{y})$. Fairness measures satisfying this property are Schur-concave.

THEOREM 3. (Schur-concavity.) A fairness measure satisfying Axioms 1-5 is Schur-concave:

$$f(\mathbf{x}) \leq f(\mathbf{y}), \quad \text{if } \mathbf{x} \preceq \mathbf{y}. \quad (9)$$

Next we present several other properties of fairness measures satisfying the axioms.

COROLLARY 2. Equal-resource allocation is fairest. A fairness measure $f(\mathbf{x})$ generated from Axioms 1-5 is maximized by equal-resource allocations, i.e.

$$f(\mathbf{1}_n) = \max_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}). \quad (10)$$

COROLLARY 3. Collecting fixed-tax is unfair. If an fixed amount $c > 0$ of the resource is subtracted from each user (i.e. $x_i - c$ for all i), the resulting fairness measure decreases

$$f(\mathbf{x} - c \cdot \mathbf{1}_n) \leq f(\mathbf{x}), \quad \forall c > 0, \quad (11)$$

where $c > 0$ must be small enough such that all elements of $\mathbf{x} - c \cdot \mathbf{1}_n$ are positive.

COROLLARY 4. Inactive user achieves no fairness. Removing users with zero resources does not change fairness:

$$f(\mathbf{x}, \mathbf{0}_n) = f(\mathbf{x}), \quad \forall n \geq 1, \quad (12)$$

4. A FAMILY OF FAIRNESS MEASURES

For any given function $g(y)$ satisfying the condition in Axiom 5, we can generate a unique expression for $f(\mathbf{x})$. Such a $f(\mathbf{x})$ is a well-defined fairness measure if it satisfies Axioms 1-5. We then refer to the corresponding $g(y)$ as a generator of the fairness measure. Unfortunately, it is difficult to find the entire set of proper generators $g(y)$, since we can not derive a necessary condition for the Axiom of Irrelevance of Partition to hold in general. Among the forms of $g(y)$ functions we have tried so far (e.g., logarithm, polynomial, exponential, and their combinations), all the resulting fairness measures collapse to the family generated by power functions.

In this section, we consider power mean $g(y) = y^\beta$ parameterized by β and derive the resulting family of fairness measures, which indeed satisfy all the axioms. From here on, we replace Equation (3) in Axiom 5 by

$$f(\mathbf{x}^1, \mathbf{x}^2) = f(w(\mathbf{x}^1), w(\mathbf{x}^2)) \cdot \left(\sum_{i=1}^2 s_i \cdot f^\beta(\mathbf{x}^i) \right)^{\frac{1}{\beta}},$$

where the weights s_i are given by Equation (5).

THEOREM 4. (Fairness measures for power mean.) For power mean $g(y) = y^\beta$ with parameter β , Axioms 1-5 define a unique family of fairness measures as follows

$$f(\mathbf{x}) = \left[\sum_{i=1}^n \left(\frac{x_i}{\sum_j x_j} \right)^{1-\beta r} \right]^{\frac{1}{\beta}}, \quad \text{for } \beta r \leq 1. \quad (13)$$

$$f(\mathbf{x}) = - \left[\sum_{i=1}^n \left(\frac{x_i}{\sum_j x_j} \right)^{1-\beta r} \right]^{\frac{1}{\beta}}, \quad \text{for } \beta r \geq 1. \quad (14)$$

where $r = \frac{1-\rho}{\beta}$ is a constant exponent, which determines the growth rate of maximum fairness as population size n increases, i.e.

$$f(\mathbf{1}_n) = n^r \cdot f(\mathbf{1}). \quad (15)$$

For different parameter β , the fairness measures given in (21) and (14) above are equivalent up to a constant exponent r . Let $f_{\beta,r}$ denote the fairness measure with parameters β and r , we have

$$f_{\beta,r}(\mathbf{x}) = f_{\beta,r,1}^r(\mathbf{x}). \quad (16)$$

According to theorem 1, since r determines the growth rate of maximum fairness as population size n increases, without

loss of generality, we choose $r = 1$ such that the maximum average fairness per user is a constant $\frac{f(\mathbf{1}_n)}{n} = f(1)$. Therefore, from a user's perspective, her individual perception of maximum fairness is independent of the population size of the system. We can use a unified representation of the proposed fairness measurers as follows

$$f(\mathbf{x}) = \text{sign}(1 - \beta) \cdot \left[\sum_{i=1}^n \left(\frac{x_i}{\sum_j x_j} \right)^{1-\beta} \right]^{\frac{1}{\beta}}. \quad (17)$$

where $\text{sign}(\cdot)$ is the sign function.

The family of fairness measures in (17) recovers many existing form of fairness indexes and utility functions. For example, when $\beta = 1$ (i.e. harmonic mean is used in Axiom 5), we get Jain's index $J(\mathbf{x}) = f(\mathbf{x})/n$. For $0 < \beta < 1$ and $\beta > 1$, we obtain α -fair utility functions with $\alpha = \beta$. Using the notation $w(\mathbf{x}) = \sum_j x_j$, we summarize the important special cases in Table 4. We use $H(\cdot)$ to denotes the Shannon entropy function. The fairness indexes and utility functions in Table 4 are now connected under our unifying framework.

β	Mean	Fairness Measure	Name
∞	maximum	$-\max_i \left\{ \frac{\sum_i x_i}{x_i} \right\}$	Max ratio
$(1, \infty)$		$-\left[(1 - \beta) U_{\alpha=\beta} \left(\frac{\mathbf{x}}{w(\mathbf{x})} \right) \right]^{\frac{1}{\beta}}$	α -fair utility ¹
$(0, 1)$		$\left[(1 - \beta) U_{\alpha=\beta} \left(\frac{\mathbf{x}}{w(\mathbf{x})} \right) \right]^{\frac{1}{\beta}}$	α -fair utility ¹
0	geometric	$e^{H\left(\frac{\mathbf{x}}{w(\mathbf{x})}\right)}$	Entropy
$(0, -1)$		$\left[\sum_{i=1}^n \left(\frac{x_i}{w(\mathbf{x})} \right)^{1-\beta r} \right]^{\frac{1}{\beta}}$	No name
-1	harmonic	$\frac{(\sum_i x_i)^2}{\sum_i x_i^2} = n \cdot J(\mathbf{x})$	Jain's index
$(-1, -\infty)$		$\left[\sum_{i=1}^n \left(\frac{x_i}{w(\mathbf{x})} \right)^{1-\beta r} \right]^{\frac{1}{\beta}}$	No name
$-\infty$	minimum	$\min_i \left\{ \frac{\sum_i x_i}{x_i} \right\}$	Min ratio

Table 2: Recover previous results as special cases of our axiomatic construction.

There have been two different approaches for characterizing fairness. The index-based approach in [1, 2, 3, 4, 5, 6] gives a direct measure of fairness by mapping an arbitrary resource allocation vector to a real number, while the the utility-based approach in [7, 8, 10, 11, 12] quantifies fairness for maximizers of the α -fair utility function. The axiomatic construction of fairness in this paper provide a unified framework for the two approaches. In Table 4, the family of fairness measures recovers the utility-based approach for $\beta > 0$, and the index-based approach for $\beta < 0$.

5. ENGINEERING IMPLICATIONS

¹This is the fairness component of α -fair utility as shown in Section 6.1. In addition, since the α -fair utility is discontinuous at $\alpha = 1$, the proportional-fair utility can be recovered by making taking the limit of our fairness measure as $\beta \rightarrow 1$.

5.1 Additional Properties

In addition to the properties proven in Corollary 1-3 in Section III, the fairness measures in (17) generated by the power mean $g(y) = y^\beta$ allows the following properties to be proved.

COROLLARY 5. *Counting the number of inactive users. The fairness measures in (17) quantify the number of inactive users in the system. When $f < 0$ is negative, $f(\mathbf{x}) \rightarrow -\infty$ if any user is assigned zero resource. When $f > 0$ is positive, for a given allocation \mathbf{x} ,*

$$\# \text{ of users with zero rate} \leq n - f(\mathbf{x}) \quad (18)$$

$$\max \text{ user rate} \geq \frac{\sum_i x_i}{f(\mathbf{x})} \quad (19)$$

COROLLARY 6. *Inequality condition for poor and rich. If we vary resource allocation of user i by a small amount ϵ , while not affecting other users' allocation, the fairness measures in (17) increases if and only if $x_i < x_f = \left(\frac{\sum_j x_j}{\sum_j x_j^{1-\beta}} \right)^{\frac{1}{\beta}}$ and $0 < \epsilon < x_f - x_i$, i.e.,*

$$f(x_1, \dots, x_i, \dots, x_n) < f(x_1, \dots, x_i + \epsilon, \dots, x_n). \quad (20)$$

This corollary implies that x_f serves as a threshold for identifying poor and rich users, since assigning an additional ϵ amount of resource to user i improves fairness if $x_i < x_f$, while the same assignment reduces fairness if $x_i > x_f$.

COROLLARY 7. *Lower bound under box-constraints. If a resource allocation $\mathbf{x} = [x_1, x_2, \dots, x_n]$ satisfies box-constraints, i.e., $x_{\min} \leq x_i \leq x_{\max}$ for all i , the fairness measures in (17) is lower bounded by*

$$f(\mathbf{x}) \geq \text{sign}(1 - \beta) \cdot \frac{(\mu \Gamma^{1-\beta} + 1 - \mu)^{\frac{1}{\beta}}}{(\mu \Gamma + 1 - \mu)^{\frac{1}{\beta} - 1}}. \quad (21)$$

where $\Gamma = \frac{x_{\max}}{x_{\min}}$ and $\mu = \frac{\Gamma - \Gamma^{1-\beta} - \beta(\Gamma - 1)}{\beta(\Gamma - 1)(\Gamma^{1-\beta} - 1)}$. The lower bound is active when a μ fraction of users receive x_{\max} and the remaining $1 - \mu$ fraction of users receive x_{\min} .

5.2 Choosing β

Parameter β determines the choice of generator function $g(y) = y^\beta$ in Axiom 5. In this section, we investigate engineering meanings of β , and give examples on how to choose proper β for different applications. We use $f_\beta(\mathbf{x})$ to denote the fairness measures in (17), parameterized by β . we first prove that for a given resource allocation \mathbf{x} , fairness $f_\beta(\mathbf{x})$ is monotonic as $\beta \rightarrow 1$. Its engineering implication is discussed next.

THEOREM 5. *(Monotonicity with respect to β .) The fairness measures in (17) is negative and monotonically decreasing for $\beta \in (1, \infty)$. For $\beta \in (-\infty, 1)$, it is positive and monotonically increasing:*

$$\frac{\partial f_\beta(\mathbf{x})}{\partial \beta} \leq 0 \text{ for } \beta \in (1, \infty), \quad (22)$$

$$\frac{\partial f_\beta(\mathbf{x})}{\partial \beta} \geq 0 \text{ for } \beta \in (-\infty, 1). \quad (23)$$

As $\beta \rightarrow 1$, f point-wise converges to constant values:

$$\lim_{\beta \uparrow 1} f_\beta(\mathbf{x}) = n \text{ and } \lim_{\beta \downarrow 1} f_\beta(\mathbf{x}) = -n. \quad (24)$$

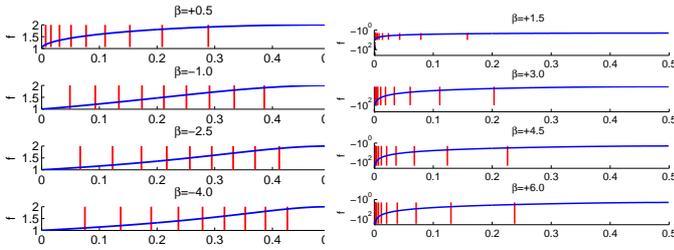


Figure 1: Plot of the fairness measure $f_\beta(\theta, 1 - \theta)$ for resource allocation $\mathbf{x} = [\theta, 1 - \theta]$ and different choices of $\beta = \{-4.0, -2.5, -1.0, +0.5\}$ and $\beta = \{+1.5, +3.0, +4.5, +6.0\}$, respectively. It can be observed that $f_\beta(\theta, 1 - \theta)$ is monotonic as $\beta \rightarrow 1$. Further, smaller values of $|1 - \beta|$ results in a steeper incline over small θ , i.e., the low-fairness region.

The monotonicity of fairness measures $f_\beta(\mathbf{x})$ on $\beta \in (-\infty, 1)$ and $\beta \in (1, \infty)$ gives an engineering interpretation of parameter β . Figure 1 plots fairness $f_\beta(\theta, 1 - \theta)$ for resource allocation $\mathbf{x} = [\theta, 1 - \theta]$ and different choices of $\beta = \{-4.0, -2.5, -1.0, +0.5\}$ and $\beta = \{+1.5, +3.0, +4.5, +6.0\}$. The vertical bars in the figure represent the level sets of function f , for values $f_\beta(\theta_i, 1 - \theta_i) = \frac{i}{10}(f_{max} - f_{min})$ for $i = 1, 2, \dots, 9$. For fixed resource allocations, since f increases as β approaches 1, it is observed that the level sets of f are pushed toward the region with small θ (i.e., the low-fairness region), resulting in a steeper incline in the region. In the extreme case of $\beta = 1$, all level sets align with the y-axis in the plot. The fairness measure f point-wise converges to step functions $f_\beta(\theta, 1 - \theta) = 2$ and $f_\beta(\theta, 1 - \theta) = 2$, respectively. Therefore, parameter β characterized the shape of the fairness measures: a smaller value of $|1 - \beta|$ (i.e., β is closer to 1) causes the level sets to be more condensed in the low-fairness region.

By varying parameter β , the proposed fairness measure f can be tuned for different applications. If the fairness measure f is used for classifying different resource allocations, a larger β is desirable, since it gives more quantization levels in low-fairness region and provides finer granularity control for unfair resource. On the other hand, if the fairness measure f is used as an objective function, a smaller β is desirable, since it has a steeper incline in the low-fairness region and give more incentive for the system to operate in the high-fairness region.

6. UNDERSTANDING α FAIRNESS

Due to Axiom 2, the Axiom of Homogeneity, our fairness measures only express desirability over the $n - 1$ -dimension subspace orthogonal to the $\mathbf{1}_n$ vector. The component of vectors along the vector $\mathbf{1}_n$ describes another quantity used to classify the efficiency of an allocation, as a function of the sum of resources $w(\mathbf{x})$.

We focus in this section on α -fair utility function:

$$\sum_i U_\alpha(x_i), \text{ where } U_\alpha(x) = \begin{cases} \frac{x^{1-\alpha}}{1-\alpha} & \alpha \geq 0, \alpha \neq 1 \\ \log(x) & \alpha = 1 \end{cases}. \quad (25)$$

We first show that the α -fairness network utility function can be factored into two components: one corresponding to the family of fairness measures we constructed and the other corresponding to efficiency. We then demonstrate that, for a fixed α , the factorization can be viewed as a single point on

the optimal tradeoff curve between fairness and efficiency. Furthermore, this particular point is one where maximum emphasis is placed on fairness while maintaining Pareto optimality of the allocation. This allows us to quantitatively interpret the belief of ‘‘larger α more fair?’’ across all $\alpha \geq 0$.

6.1 Factorization of α -fair Utility Function

Re-arranging the terms of the equation in Table 4, we have

$$\begin{aligned} U_{\alpha=\beta}(\mathbf{x}) &= \frac{1}{1-\beta} |f_\beta(\mathbf{x})|^\beta \left(\sum_i x_i \right)^{1-\beta} \\ &= |f_\beta(\mathbf{x})|^\beta \cdot U_\beta \left(\sum_i x_i \right), \end{aligned} \quad (26)$$

where $U_\beta(\sum_i x_i)$ is the one-dimensional version of the α -fair utility function with $\alpha = \beta$. For $\beta \rightarrow 1$, it is easy to show that our fairness measure $f_\beta(\mathbf{x})$, multiplied by a function of throughput $\sum_i x_i$, equals α -fair utility function with $\alpha = 1$. Similarly, for $\beta \rightarrow \infty$, it equals α -fair utility function as $\alpha \rightarrow \infty$. Therefore, Equation (26) also holds for proportional fairness at $\alpha = 1$ and max-min fairness at $\alpha \rightarrow \infty$.

Equation (26) demonstrates that the α -fair utility functions can be factorized as the product of two components: a fairness measure, $|f_\beta(\mathbf{x})|^\beta$, and an efficiency measure, $U_\beta(\sum_i x_i)$.

The fairness measure $|f_\beta(\mathbf{x})|^\beta$ only depends on the normalized distribution, $\mathbf{x}/(\sum_i x_i)$, of resources (due to Axiom 2), while the efficiency measure is a function of the sum resource $\sum_i x_i$.

The factorization of α -fair utility functions decouples the fairness component from the efficiency component to tackle issues such as fairness-efficiency tradeoff and feasibility of \mathbf{x} under a given constraint set of utility maximization. For example, it helps to explain the counter-intuitive throughput behavior in [12]: an allocation vector that maximizes the α -fair utility with a larger α may not be less efficient, because the α -fair utility incorporates both fairness and efficiency at the same time.

6.2 Pareto Optimality in Fairness-Efficiency Tradeoffs

Although Corollary 2 states equal allocation is fairest, an α -fair allocation may not have an equal distribution. This is because the additional efficiency component of sum-throughput in (26) can skew desirability away from an equal distribution. The magnitude of this skewing depends on the fairness parameter ($\alpha = \beta$), the constraint set of \mathbf{x} , and the relative importance of fairness and efficiency.

Guided by the product form of (26), we consider a scalarization of maximization of the two objectives, fairness and efficiency:

$$\Phi_\lambda(\mathbf{x}) = \lambda \ell(f_\beta(\mathbf{x})) + \ell \left(\sum_i x_i \right), \quad (27)$$

where $\beta \in (0, 1) \cup (1, \infty)$ is fixed, $\lambda \in [0, \infty)$ absorbs the exponent β in the fairness component of (26) and is a weight specifying the relative emphasis placed on the fairness, and

$$\ell(y) = \text{sign}(y) \log(|y|). \quad (28)$$

The use of the log function later recovers the product in the factorization of (26) from the sum in the scalarized (27).

An allocation vector \mathbf{x} is Pareto dominated by \mathbf{y} if $x_i \leq y_i$ for all i and $x_i < y_i$ for at least some i . An allocation is

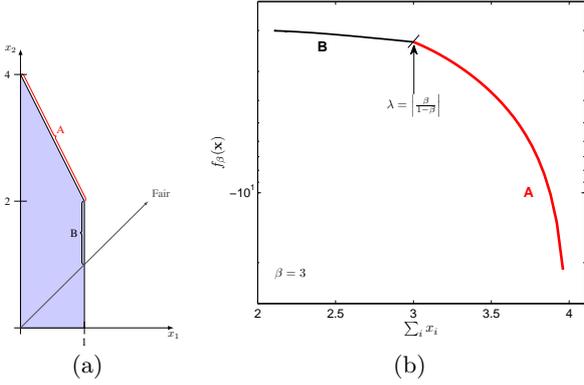


Figure 2: (a) feasible region where over emphasis of fairness violates Pareto dominance, and (b) its fairness-efficiency tradeoff for $\beta = 3$. The Region A corresponds to Pareto optimal solutions. The Region B is when the condition of Theorem 6 is violated, and solutions are more fair, but no longer Pareto optimal.

Pareto optimal if it is not Pareto dominated by any other feasible allocation. If the relative emphasis on efficiency is sufficiently high, Pareto optimality of the solution can be maintained. To preserve Pareto optimality, we require that if \mathbf{y} Pareto dominates \mathbf{x} , then $\Phi_\lambda(\mathbf{y}) > \Phi_\lambda(\mathbf{x})$.

THEOREM 6. *Preserving Pareto optimality.* Let $\beta \in [0, \infty) \setminus \{1\}$ be fixed and (27) be the scalarization of the fairness index and sum throughput that is used as the objective for maximization. If

$$\lambda \leq \left| \frac{\beta}{1-\beta} \right|, \quad (29)$$

then Pareto dominant allocations are preserved.

Consider the set of maximizers of (27) for λ in the range in Theorem 6:

$$\mathbb{P} = \left\{ \mathbf{x} : \mathbf{x} = \arg \max_{\mathbf{x} \in \mathbb{R}} \Phi_\lambda(\mathbf{x}), \forall \lambda \leq \left| \frac{\beta}{1-\beta} \right| \right\}. \quad (30)$$

When weight $\lambda = 0$, the corresponding points in \mathbb{P} is throughput optimal. When weight $\lambda = \left| \frac{\beta}{1-\beta} \right|$, it can be shown that the factorization in (26) is equivalent to (27). Therefore, α -fairness corresponds to the solution of an optimization which places the maximum emphasis on the fairness measure parameterized by $\beta = \alpha$ while preserving Pareto optimality. Allocations in \mathbb{P} corresponding to other values of λ achieve a tradeoff between fairness and efficiency, while Pareto optimality is preserved.

Figure 2(b) illustrates an optimal fairness-efficiency tradeoff curve $\left\{ [f_\beta(\mathbf{x}), \sum_i x_i], \forall \mathbf{x} = \arg \max_{\mathbf{x} \in \mathbb{R}} \Phi_\lambda(\mathbf{x}), \forall \lambda \right\}$ corresponding to the constraint set shown in Figure 2(a). The set of optimizers \mathbb{P} in (30), which is obtained by maximizing Pareto optimal utilities (27), is shown by curve A in Figure 2(b).

6.3 Why Larger α is More Fair

In the previous section we demonstrated the factorization (26) is an extreme point on the tradeoff curve between fairness and efficiency for fixed $\beta = \alpha$. What happens when α becomes bigger?

We denote by $\nabla_{\mathbf{x}}$ the gradient operator with respect to the vector \mathbf{x} . For a differentiable function, we use the standard inner product ($\langle \mathbf{x}, \mathbf{y} \rangle = \sum_i x_i y_i$) between the gradient of the function and a normalized vector to denote the directional derivative of the function.

THEOREM 7. *(Monotonicity of fairness-efficiency reward ratio.)* Let allocation \mathbf{x} be given. Define $\boldsymbol{\eta} = \frac{1}{n} \mathbf{1}_n - \frac{\mathbf{x}}{\sum_i x_i}$ as the vector pointing from the allocation to the nearest fairness maximizing solution. Then the fairness-efficiency reward ratio

$$\frac{\left\langle \nabla_{\mathbf{x}} U_{\alpha=\beta}(\mathbf{x}), \frac{\boldsymbol{\eta}}{\|\boldsymbol{\eta}\|} \right\rangle}{\left\langle \nabla_{\mathbf{x}} U_{\alpha=\beta}(\mathbf{x}), \frac{\mathbf{1}_n}{\|\mathbf{1}_n\|} \right\rangle}, \quad (31)$$

is non-decreasing with α , i.e. higher α gives a greater relative reward for fairer solutions.

The choice of direction $\boldsymbol{\eta}$ is a direct result of Axiom 2 and Corollary 2, which together imply $\boldsymbol{\eta}$ is the direction that most increases fairness and is orthogonal to increases in efficiency.

An increase in either fairness or efficiency is a “desirable” outcome. The choice of α dictates exactly how desirable one objective is relative to the other (from a fixed allocation). Theorem 7 states that, with a larger α , there is a larger component of the utility function gradient in the direction of fairer solutions, relative to the component in the direction of more efficiency. Notice however that comparison must be in terms of the ratio between these two gradient components rather than the magnitude of the gradient, and both fairness and efficiency may increase simultaneously.

This result provides a justification for the belief that larger α is “more fair”, for any $\alpha \in [0, \infty)$. Figure 3 depicts how this ratio increases with $\alpha = \beta$ for some examples allocations.

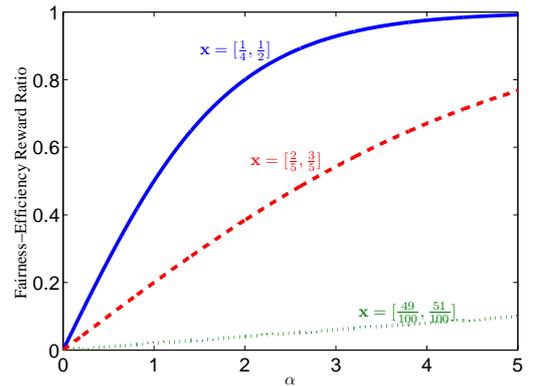


Figure 3: Monotonic behavior of the ratio between the reward for increased fairness and reward for increased efficiency as a function of $\alpha = \beta$. Three fixed allocations are considered, and solutions that are already more fair have a lower ratio. However the ratio still increases monotonically.

7. REFERENCES

- [1] M.A. Marson and M. Gerla, "Fairness In Local Computing Networks," in *Proceeding of IEEE ICC*, 1982.
- [2] J.W. Wong and S.S. Lam, "A Study Of Fairness In Packet Switching Networks," in *Proceedings of IEEE ICC*, 1982.
- [3] R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer system," *DEC Technical Report 301*, 1984.
- [4] M. Dianati, X. Shen, and S. Naik, "A New Fairness Index for Radio Resource Allocation in Wireless Networks," in *Proceedings of WCNC*, 2005.
- [5] C. E. Koksal, H. I. Kassab, and H. Balakrishnan, "An analysis of short-term fairness in wireless media access protocols," in *Proceedings of ACM SIGMETRICS*, 2000.
- [6] M. Bredel and M. Fidler, "Understanding Fairness and its Impact on Quality of Service in IEEE 802.11," in *Proceedings of IEEE INFOCOM*, 2009.
- [7] F. P. Kelly, A. Maulloo and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237-252, 1998.
- [8] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Transactions Networking*, vol. 8, no. 5, pp. 556-567, Oct. 2000.
- [9] M. Uchida and J. Kurose, "An Information-Theoretic Characterization of Weighted α -Proportional Fairness," in *Proceedings of IEEE INFOCOM*, 2009.
- [10] T. Bonald and L. Massoulié, "Impact of Fairness on Internet Performance," in *Proceedings of ACM Sigmetrics*, 2001.
- [11] L. Massoulié and J. Roberts, "Bandwidth sharing: objectives and algorithms," *IEEE/ACM Transactions Networking*, vol. 10, no. 3, pp. 320-328, Jun. 2002.
- [12] A. Tang, J. Wang and S. Low, "Counter-intuitive Behaviors in Networks under End-to-end Control," *IEEE /ACM Transactions on Networking*, vol. 14, no. 2, pp. 355-368, April 2006.
- [13] R. Bhargava, A. Goelt, and A. Meyerson, "Using Approximate Majorization to Characterize Protocol Fairness," in *Proceedings of ACM Sigmetrics*, 2001
- [14] A. W. Marshall and I. Olkin, "Inequalities: Theory of Majorization and its applications," *Academic Press*, 1979.
- [15] A. Kolmogoroff, "Sur la notion de la moyenne", *Atti della R. Accademia nazionale dei Lincei*, volumn 12, 1930.
- [16] A. Renyi, "On measures of information and entropy", in *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability*, 1960.
- [17] P. Erdos, "On the distribution function of additive functions", *Ann. of Math.*, Vol. 47 pp. 1-20, 1946.
- [18] J.F. Nash, "The Bargaining Problem", *Econometrica*, vol. 18, no. 2, pp. 155-162, 1950.
- [19] L.S. Shapley, "A Value for n-person Games", *Contributions to the Theory of Games*, vol. II, pp. 307-317, Princeton University Press, 1953.
- [20] S. Boyd and L. Vandenberghe "Convex optimization", *Cambridge University Press*, 2005.
- [21] T. Lan, D. Kao, M. Chiang, and A. Sabharwal "An Axiomatic Approach to Fairness", *Technical report*, available at www.princeton.edu/~tlan/fairness.pdf, 2009.